

LISTING OF CLAIMS:

This listing of claims will replace all prior versions, and listings, of claims in the application:

Claims 1 – 18 (Canceled)

19. (Withdrawn) A method of representing an audio signal for machine learning comprising:

(a) creating a perceptual representation of said audio signal by performing a frequency domain transform on at least one time-sampled window of a digital representation of said audio signal, said perceptual representation comprising component magnitudes of constituent frequency vectors that comprise said audio signal;

(b) calculating a magnitude of each constituent frequency vector within said audio signal;

(c) grouping each of said constituent frequency vectors into a number of frequency bands;

(d) calculating an average magnitude of said constituent frequency vectors within each of said frequency bands; and

(e) arranging said magnitudes into a learning representation.

20. (Withdrawn) The method according to claim 19 wherein said frequency domain transform is a Fast Fourier Transform.

21. (Withdrawn) The method according to claim 19 wherein an average magnitude of said constituent frequency vectors within each of said frequency bands further comprises an aggregate average magnitude over a plurality of said time-sampled windows.

22. (Withdrawn) The method according to claim 21 where said plurality of time-sampled windows comprises 12 time-sampled windows.
23. (Withdrawn) The method according to claim 19 wherein no said frequency band includes any frequency greater than 11 kHz.
24. (Withdrawn) The method according to claim 19 wherein said frequency bands grow in size according to the golden ratio of frequency with respect to pitch.
25. (Withdrawn) The method according to claim 19 further comprising the step of converting said audio signal into a pulse code modulated bitstream for processing by said frequency domain transform.
26. (Withdrawn) A computer readable storage medium, storing therein a program of instructions for causing a computer to execute process of representing an audio signal for machine learning, said process comprising the steps of:
- (a) creating a perceptual representation of said audio signal by performing a frequency domain transform on at least one time-sampled window of a digital representation of said audio signal, said perceptual representation comprising component magnitudes of constituent frequency vectors that comprise said audio signal;
 - (b) calculating a magnitude of each constituent frequency vector within said audio signal;
 - (c) grouping each of said constituent frequency vectors into a number of frequency bands;
 - (d) calculating an average magnitude of said constituent frequency vectors within each of said frequency bands; and

- (e) arranging said magnitudes into a learning representation.

Claims 27 – 43 (Canceled)

44. (Withdrawn) An apparatus for representing an audio signal for machine learning comprising:

- (a) a means for performing a frequency domain transform on at least one time-sampled window of a digital representation of said audio signal, said perceptual representation comprising component magnitudes of constituent frequency vectors that comprise said audio signal;
- (b) a means for calculating a magnitude of each constituent frequency vector;
- (c) a means for grouping each of said constituent frequency vectors into a number of frequency bands;
- (d) a means for calculating an average magnitude of said constituent frequency vectors within each of said frequency bands; and
- (e) a means for arranging said magnitudes into a learning representation.

45. (Withdrawn) The apparatus according to claim 44 wherein said means for performing a frequency domain transform comprises a means for performing a Fast Fourier Transform.

46. (Withdrawn) The apparatus according to claim 44 wherein no said frequency band includes any frequency greater than 11 kHz.

47. (Withdrawn) The apparatus according to claim 44 wherein said frequency bands grow in size according to the golden ratio of frequency with respect to pitch.

48. (Withdrawn) The apparatus according to claim 44 further comprising a means for converting said audio signal into a pulse code modulated bitstream for processing by said frequency domain transform.

49. (Currently Amended) A method of extracting classifying data from an audio signal, the method comprising the steps of:

~~processing~~ transforming a perceptual representation of the audio signal into a learning representation of the audio signal; [[and]]

~~inputting~~ transmitting the learning representation to ~~into~~ a multi-stage classifier, the multi-stage classifier comprising:

a first stage [[of]] having a plurality of support vector machine classifiers, each support vector machine classifier trained to identify one out of a plurality of audio classification categories and generate a metalearner vector value reflecting how closely the audio signal conforms to the one out of the plurality of audio classification categories, and

a final stage having a metalearner classifier, the metalearner classifier using the generated metalearner vector ~~each support vector machine classifier trained to identify one out of a plurality of audio classification categories and~~

~~where the support vector machine classifiers are used to generate a metalearner vector that allows the final stage metalearner classifier to classify the audio signal into one out of the plurality of audio classification categories; and~~

generating classification category information for the audio signal based on results produced by the metalearner classifier.

~~, each support vector machine classifier outputting a value reflecting how closely the audio signal conforms to the one out of the plurality of audio classification categories, each value then used in the metalearner vector.~~

50. (Previously Presented) The method of claim 49 wherein the final stage metalearner classifier is a neural network classifier.

51. (Canceled)

52. (Previously Presented) The method of claim 49 wherein said audio classification categories comprises classifications by musical artist.

53. (Previously Presented) The method of claim 49 wherein the learning representation comprises dividing the perceptual representation of the audio signal into a plurality of time slices.

54. (Previously Presented) The method of claim 49 wherein the learning representation comprises dividing the perceptual representation of the audio signal into a plurality of frequency bands.

55. (Previously Presented) A computer readable storage medium, storing therein a program of instructions for causing a computer to execute a process of extracting classifying data from an audio signal, the process comprising the steps of:

processing a perceptual representation of the audio signal into a learning representation of the audio signal; and

inputting the learning representation into a multi-stage classifier, the multi-stage classifier comprising a first stage of support vector machine classifiers and a final stage metalearner classifier, each support vector machine classifier trained to identify one out of a plurality of audio classification categories and where the support vector machine classifiers are used to generate a metalearner vector that allows the final stage metalearner classifier to classify the audio signal into one out of the plurality of audio classification categories, each support vector machine classifier outputting a value reflecting how closely the audio signal conforms to the one out of the plurality of audio classification categories, each value then used in the metalearner vector.

56. (Previously Presented) The computer readable storage medium of claim 55 wherein the final stage metalearner classifier is a neural network classifier.

57. (Canceled)

58. (Previously Presented) The computer readable storage medium of claim 55 wherein said audio classification categories comprises classifications by musical artist.

59. (Previously Presented) The computer readable storage medium of claim 55 wherein the learning representation comprises dividing the perceptual representation of the audio signal into a plurality of time slices.

60. (Previously Presented) The computer readable storage medium of claim 55 wherein the learning representation comprises dividing the perceptual representation of the audio signal into a plurality of frequency bands.

61. (Previously Presented) An apparatus for classifying an audio signal comprising:

means for processing a perceptual representation of the audio signal into a learning representation of the audio signal; and

a multi-stage classifier, the multi-stage classifier further comprising a first stage of support vector machine classifiers and a final stage metalearner classifier, each support vector machine classifier trained to identify one out of a plurality of audio classification categories from the learning representation of the audio signal and where the support vector machine classifiers are used to generate a metalearner vector that allows the final stage metalearner classifier to classify the audio signal into one out of the plurality of audio classification categories, each support vector machine classifier outputting a value reflecting how closely the audio signal conforms to the one out of the plurality of audio classification categories, each value then used in the metalearner vector.

62. (Previously Presented) The apparatus of claim 61 wherein the final stage metalearner classifier is a neural network classifier.

63. (Canceled)

64. (Previously Presented) The apparatus of claim 61 wherein said audio classification categories comprises classifications by musical artist.

65. (Previously Presented) The apparatus of claim 61 wherein the learning representation comprises dividing the perceptual representation of the audio signal into a plurality of time slices.

66. (Previously Presented) The apparatus of claim 61 wherein the learning representation comprises dividing the perceptual representation of the audio signal into a plurality of frequency bands.